

UNITED STATES PATENT APPLICATION

FOR

ADAPTIVE HYBRID FOR A
DIGITAL SUBSCRIBER LINE ANALOG FRONT END

INVENTOR:

KIM FUNG

Prepared by:

WILLIAM D. DAVIS
DAVIS & JOHNSON, L.L.P.
P.O. BOX 90698
AUSTIN, TEXAS 78709

Attorney Docket No. 75622.P0035

EL 566102397 US

"Express Mail" mailing label number

May 1, 2001

Date of Deposit

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

William D. Davis

(Typed or printed name of person mailing paper or fee)

William D. Davis

(Signature of person mailing paper or fee)

ADAPTIVE HYBRID FOR A
DIGITAL SUBSCRIBER LINE ANALOG FRONT END

FIELD OF THE INVENTION

This invention relates to the field of telecommunications. In 5 particular, this invention is drawn to combining digital and analog telecommunications functions.

BACKGROUND OF THE INVENTION

Numerous communication protocol standards have developed to enable using the pre-existing plain old telephone system (POTS) 10 infrastructure for carrying digital data. Although the public switched telephone network (PSTN) is digital in nature, the connection between subscribers and the central office which serves as an entry point to the PSTN is analog. As a result, modems are used for bi-directional communication of digital data on the analog channel between the subscriber and the central 15 office. Modems convert the communicated information between the digital and analog domains in accordance with the particular communication protocol.

Some communication protocols are designed to rely on the voiceband region of the analog channel to convey information. As a result, when a 20 subscriber line is in use by such a voiceband modem, the line is unavailable for simultaneous voice communications.

Digital subscriber lines (xDSL) services can provide significantly higher data transmission rates by utilizing communication bandwidth beyond but excluding the voiceband. As a result, xDSL services may simultaneously co- 25 exist with voiceband communications.

Modems or other devices designed for communicating digital data on the analog channel utilize an analog front end for transmitting as well as receiving information from the subscriber line. The analog front end conditions signals communicated to or from the subscriber line before

5 providing the conditioned signal to the subscriber line for transmission or to a digital signal processor for interpretation.

A hybrid circuit, for example, is used to address echoes resulting from using the same two wires for both transmission and reception on the analog channel. One disadvantage of typical hybrid designs is that the high order

10 introduces distortion and noise into the system.

Preferably, a modem has the ability to support the higher data rates of xDSL when available. Due to the geographical limitations on xDSL, voiceband modems are still needed to ensure a reliable means of communication. One risk averse solution implements the functionality of

15 voiceband and xDSL modems through the use of a chipset mounted on a common circuit board.

The combination of the circuitry onto a common board introduces new problems. In particular, the voiceband and xDSL modems rely on clocks of different frequencies which are not multiples of each other. The interaction

20 between two clocks, for example, may result in intermodulation, synchronization problems, or other issues which interfere with the digital signal processor's ability to properly interpret received information.

SUMMARY OF THE INVENTION

Methods and apparatus incorporated into an analog front end for communicating with a subscriber line enable a high degree of integration.

Upstream data to be transmitted to a subscriber line is pre-processed to 5 eliminate even images. A power spectral density shaping filter subsequently substantially eliminates undesired energy in the upstream data signal. A tunable hybrid network is coupled to the subscriber line and the transmit block to extract downstream data from the subscriber line. A high pass filter further rejects upstream data from the downstream data signal. The power 10 spectral density shaping filter and the high pass filter enable the use of a tuned low order hybrid network to achieve adequate near end echo cancellation. The analog front end may include an additional analog channel to enable voiceband (e.g., v.90) communication concurrent with non-voiceband (e.g., xDSL) operation. Sample rate conversion is utilized to 15 avoid the use of multiple independent clocks for otherwise incompatible clocking requirements.

One embodiment of an analog front end apparatus includes a transmit block coupled to transmit discrete multitone modulated upstream data to a subscriber line. A hybrid network is coupled to the subscriber line and the 20 transmit block. A receive block is coupled to the hybrid for receiving discrete multitone modulated downstream data from the subscriber line. The transmit block, hybrid network, and receive block reside within a same integrated circuit package.

The hybrid network includes a hybrid input, a receive input, and a 25 hybrid output. The receive input is capacitively coupled to a subscriber line

carrying the upstream and downstream data signals. The hybrid input is
capacitively coupled to receive the upstream data signal from the transmit
block driver. The driver also provides the upstream data signal to the
subscriber line. The hybrid output provides the extracted downstream data
5 signal from the subscriber line. The hybrid network may be fabricated on an
integrated circuit die. In one embodiment, the hybrid network is a
complementary metal oxide semiconductor integrated circuit.

Other features and advantages of the present invention will be
apparent from the accompanying drawings and from the detailed description
10 that follows below.

Express Mail Label

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

5 Figure 1 illustrates the communication spectrum allocated for a subscriber line.

Figure 2 illustrates an analog front end for communication with a subscriber line.

10 Figure 3 illustrates the analog front end transmit block.

Figure 4 illustrates the analog front end receive block.

Figure 5 illustrates various hybrid network implementations.

Figure 6 illustrates one embodiment of a tunable hybrid.

Figure 7 illustrates a voiceband transmit path with sample rate conversion.

15 Figure 8 illustrates a voiceband receive path with sample rate conversion.

Figure 9 illustrates a table of upsampling and downsampling factors for realizing various clock rates from a universal conversion rate.

20 Figure 10 illustrates one embodiment of a finite impulse response filter.

Figure 11 illustrates one embodiment of a polyphase pointer and serial interface pulse generator.

Figure 12 illustrates one embodiment of a polyphase pointer for the receive path of Figure 8.

DETAILED DESCRIPTION

The International Telecommunication Union (ITU) has set forth a series of recommendations for subscriber line data transmission. These recommendations are directed towards communications using the voiceband portion of the communications spectrum ("V.x" recommendations) as well as communications utilizing frequency spectrum other than the voiceband portion (e.g., "xDSL" recommendations).

The V.x recommendations have evolved over time to support ever higher data rates. ITU-T Recs. V.22 bis, V.32, V.32 bis, V.34, and V.90, for example, relate to increasing bit-per-second data rates of 2400, 9600, 14400, 33600, and 56000 bit/sec. Compression standards such as V.42 bis can further increase the effective data rates. Generally, voiceband modems will use a recommended handshaking protocol to negotiate the highest possible data rate.

ITU documentation for the more recent recommendations related to voiceband communications includes "Rec. V.90 (09/98) - A digital modem and analogue modem pair for use on the Public Switched Telephone Network (PSTN) at data signalling rates of up to 56,000 bit/s downstream and up to 33,600 bit/s upstream," "Rec. V.34 (02/98) - A modem operating at data signalling rates of up to 33 600 bit/s for use on the general switched telephone network and on leased point-to-point 2-wire telephone-type circuits," "Rec. V.32 bis (02/91) - A duplex modem operating at data signalling rates of up to 14 400 bit/s for use on the general switched telephone network and on leased point-to-point 2-wire telephone-type circuits," and "Rec. V.32 (03/93) - A family of 2-wire, duplex modems operating at data signalling rates of up to

9600 bit/s for use on the general switched telephone network and on leased telephone-type circuits."

Asymmetric digital subscriber line (ADSL) communications represent one variant of xDSL communications. Exemplary ADSL specifications are set forth in "Rec. G.992.2 (06/99) Splitterless asymmetric digital subscriber line (ADSL) transceivers" (also referred to as full rate ADSL), and "Rec. G.992-1 (06/99) - Asymmetric digital subscriber line (ADSL) transceivers" (also referred to as G.LITE).

Figure 1 illustrates the communication spectrum allocation for a subscriber line. Chart 100 compares the portions of the analog channel used by voiceband modems (POTS 110) as well as xDSL modems (ADSL 130). POTS communications typically use the voiceband range of 300-4000 Hz. ADSL is in a range of approximately 25-1100 KHz. A guard band 120 separates the POTS and ADSL ranges.

There are multiple line coding variations for xDSL. Carrierless Amplitude Phase (CAP) and Discrete Multi-Tone modulation both use the fundamental techniques of quadrature amplitude modulation (QAM). CAP is a single carrier protocol where the carrier is suppressed before transmission and reconstructed at the receiving end. DMT is a multicarrier protocol.

Figure 1 illustrates DMT line coding.

DMT modulation has been established as a standard line code for ADSL communication. The available ADSL bandwidth is divided into 256 sub-channels. Each sub-channel 134 is associated with a carrier. The carriers (also referred to as tones) are spaced 4.3125 KHz apart. Each sub-channel is modulated using quadrature amplitude modulation (QAM) and can carry

0-15 bits/Hz. The actual number of bits is allocated depending upon line conditions. Thus individual sub-channels may be carrying different numbers of bits/Hz. Some sub-channels 136 might not be used at all. ADSL uses some sub-channels 134 for downstream communication and other sub-channels

5 132 for upstream communication. The upstream and downstream sub-channels may be separated by another guard band 140.

During initialization DMT measures the signal-to-noise ratio of each sub-channel to assign a data rate. Generally, greater data rates (i.e., more bits/Hz) are assigned to the lower sub-channels because signals are attenuated
10 more at higher frequencies. DMT implementations may also incorporate rate adaption to monitor the line conditions and dynamically change the data rate for sub-channels.

Figure 2 illustrates an analog front end for communicating information on an analog channel carried by a subscriber line between a subscriber and the central office. In one embodiment, the analog front end includes circuitry for handling voiceband as well as non-voiceband (i.e., xDSL) communications. The analog front end conditions signals communicated between the subscriber line 290 and digital signal processors
15 270, 280.

20 In one embodiment, the analog front end 218 (including hybrid 260) is provided on a single substrate within an integrated circuit package. Within the integrated circuit, non-voiceband communication (e.g., xDSL) is handled by transmit block 230, receive block 240, 2-4 wire block 250, and hybrid 260. The hybrid and 2-4 wire conversion functions may be combined into a
25 common hybrid network block as illustrated. In one embodiment, the analog

front end is implemented as complementary metal oxide semiconductor (CMOS) circuitry within the integrated package. The integrated circuit may further comprise additional circuitry to support voiceband communications within the same integrated circuit package. For example, sample rate 5 converter 210 and codec 212 may be incorporated onto the same substrate as the xDSL circuitry.

The high level of integration of the analog front end is enabled predominately by four features. First, in order to ease filtering, digital signal processor 270 eliminates even images of the upstream signal through the use 10 of a Fast Fourier Transform (FFT). Second, a power spectral density shaping filter is used to further reduce unwanted images. These first two features ensure compliance with ITU specified transmit power spectrum density requirements. Third, the combination of the double rate FFT and power 15 shaping enable the use of a low order hybrid rather than a higher order distortion inducing hybrid. The hybrid is tunable to permit accommodation of varying line conditions and downstream distortion. Fourth, a high pass filter following the hybrid provides additional rejection of the transmit signal in the receive path.

Digital signal processor (DSP) 270 provides information in digital form 20 to the transmit block 230 for communication on the analog channel of the subscriber line. Figure 3 illustrates transmit block 230 in further detail.

The transmit path is digital in nature from the digital signal processor until the digital-to-analog converter 360. Data from the DSP is provided to the programmable gain amplifier 310 of the transmit block. Programmable

gain amplifier 310 provides amplification in the range of 0 to -3 dB.

Attenuation may be used to adjust the power level, if desired.

The xDSL line coding requires orthogonality between carriers.

Orthogonality may be achieved through the use of DSP FFT algorithms.

5 In one embodiment, the xDSL circuitry supports an ADSL protocol with a 138 KHz (32 tones) upstream data rate. Although the Nyquist sampling frequency for a 138 KHz upstream signal is 276 KHz, DSP 270 performs a Fast Fourier Transform at twice the necessary sampling rate (i.e., 552 KHz or 128 points instead of the 64 points typically used to process the 32 tones) to

10 eliminate the even images of the upstream data and to force the spectrum content to zero for frequencies between 142 KHz and 276 KHz. Thus in one embodiment, DSP 270 is providing upstream data values to the transmit block at 552 KHz.

A high pass filter 320 is provided to reduce spectral leakage at low frequencies caused by DSP FFT algorithms. In one embodiment, high pass filter 320 is a third order Butterworth filter with a corner frequency of approximately 12 KHz and rejection greater than 28 dB for frequencies of approximately 4 KHz or less.

Interpolator 330 interpolates the filtered signal from 552 KHz to 1.104 MHz for spectral power shaping. In one embodiment, interpolator 330 is a 4th order interpolator. The transfer function for interpolator 330 is as follows:

$$h(z)_{330} = \left(\frac{1}{N} \cdot \frac{1 - z^{-N}}{1 - z^{-1}} \right)^4$$

The ratio of the interpolator frequencies determines N ($N = \frac{f_{sout}}{f_{sin}}$). Thus,

$$N = \frac{1.104 \text{ MHZ}}{552 \text{ KHZ}} = 2.$$

The interpolated signal is processed by a power spectral density (PSD) mask or shaper 340 to ensure compliance with protocol specifications. In one embodiment PSD shaper 340 incorporates a 6th order Chebyshev filter designed to attenuate the signal by approximately 40 dBm/Hz. In one 5 embodiment, PSD shaper 340 has a corner frequency of approximately 140 KHz .

Interpolator 350 interpolates the power shaped signal from 1.104 MHz to 35.328 MHz for the purpose of rejecting upstream images. In one embodiment, interpolator 350 is a 6th order interpolator. The transfer 10 function for interpolator 350 is as follows:

$$h(z)_{350} = \left(\frac{1}{N} \cdot \frac{1 - z^{-N}}{1 - z^{-1}} \right)^6.$$

In this case, $N = \frac{35.358 \text{ MHZ}}{1.104 \text{ MHZ}} = 32$.

Digital-to-analog converter (DAC) 360 generates an analog signal from the interpolated signal. The analog signal is provided to low pass filter 370. 15 Due to the lack of even images, high DAC sampling rate, and the high rejection of the overall transmit path digital filter, low pass filter 370 may be implemented as a first order filter to reject undesired upstream images. In one embodiment, low pass filter 370 has a corner frequency of approximately 250 KHz. Driver 380 drives the upstream signal onto the subscriber line 290.

20 Referring to Figure 2, receive block 240 interfaces with subscriber line 290 through hybrid 260. In one embodiment, hybrid 260 is a tunable, first order filter. The hybrid network is discussed in further detail with respect to Figures 5-6.

Figure 4 illustrates receive block 240 in greater detail. The 4 wire signal 25 carrying the transmitted signal as well as any received communications are

passed through hybrid 410 to perform echo cancellation and to convert the 4 wires to 2 wires carrying predominately the received communications. A first programmable gain amplifier (PGA 412) permits gain adjustments from -8 to 24dB. Attenuation may be necessary in some cases due to line conditions.

5 The signal is then provided to high pass filter and PGA 414 to further reduce upstream communications in the receive path. High pass filter also permits gain adjustments from approximately 6-12 dB.

In one embodiment, high pass filter 414 is a 3rd order filter with a corner frequency of approximately 140 KHz to provide approximately 20 db

10 rejection to the upstream signal. The use of the high pass filter and PGA 414 may be preferable to the use of DACs to cancel near-end signals. Traditional DAC cancellation techniques do not account for driver noise or distortion.

The filtered signal is provided to low pass filter and PGA 416. Low pass filter and PGA 416 permits gain adjustments from approximately 0-6 dB. In one embodiment, low pass filter 416 is a 2nd order filter with a corner frequency of approximately 2 MHz. The primary purpose of low pass filter and PGA 416 is to provide anti-aliasing for analog-to-digital converter (ADC) 420. The signal is passed through another programmable gain amplifier 418 to provide for additional level control before being processed by ADC 420.

20 In one embodiment, ADC is a 3rd order sigma-delta converter. Decimator 430 serves to reduce the truncation noise contributed by ADC 420 at higher frequencies due to the sigma-delta effect. The digital signal is decimated from 35.328 MHZ to 8.836 MHZ.

25 The output of decimator 430 is provided to a programmable low pass filter 440. Low pass filter provides outband rejection to prevent the outband

signal from folding back into the baseband signal. In one embodiment low pass filter 440 is a 6th order filter. After low pass filter 440, the signal is passed to decimator 450 which decimates the signal from 8.836 to 2.208 MHz (the baseband frequency). The decimated signal is processed by high pass filter 460

5 before the 16 bit 2.208 MHz data is provided to a DSP. High pass filter 460 serves to reject any remaining audio (i.e., voiceband) components in the processed signal. In one embodiment, high pass filter 460 is a 2nd order filter.

Figure 5 compares a system level view of a hybrid network (550) with another hybrid (500). The hybrid is designed in part to cancel the signal X

10 being transmitted onto the subscriber line so that the large transmit signal is not re-introduced into the receive path. The signal X transmitted by the driver 510 is distorted by the transformer and line impedance in a manner modeled as $H1(s)$ 520. The receive path sees a superposition of the desired input signal Y and $X \cdot H1(s)$. $H2(s)$ 522 is designed to introduce the same

15 distortion as $H1(s)$ 520. When combined by differential summer 530, the result is $Y + X \cdot H1(s) - X \cdot H2(s)$. If $H1(s) = H2(s)$, the received result is the desired result, i.e., Y.

One disadvantage with this approach is that the desired receive signal, Y, is usually much smaller than X. Differential summer 530 is taking the

20 difference of two larger values in order to identify a smaller value. This tends to introduce noise into the system. Another disadvantage is that $H2(s)$ is a pole/zero filter which does not block DC. Furthermore, the pole and zero are difficult to adjust independently in practical implementations. An alternative hybrid network is illustrated as 550 for comparison.

In this case $H1(s)$ 570 modifies X such that the receive path sees $Y + H1(s) \cdot X$. A graphical representation of a typical $H1(s)$ is illustrated in chart 590. The typical $H1(s)$ has a zero at $F1$ and a pole at $F2$. A high pass filter 582 with a corner frequency of $F2$ is inserted into the receive path to produce 5 $HPF2(s) \cdot (Y + H1(s) \cdot X)$. Another high pass filter 572 with a corner frequency $F1$ is applied to the transmitted signal X to produce $HPF1(s) \cdot X$. Differential summer 580 thus produces, $HPF2(s) \cdot Y + HPF2(s) \cdot H1(s) \cdot X - HPF1(s) \cdot X$.

The corner frequency of high pass filter 572 produces a pole which offsets a zero of $H1(s)$ at the same location. The product of both filters appears 10 as another high pass filter with a corner frequency of $F1$. High pass filter 570 is selected such that $HPF2(s) \cdot H1(s) = HPF1(s)$. Thus the hybrid is able to reject the transmitted signal, X .

One advantage of this approach is that filters 572 and 582 are the same type of filter. Moreover, filters 572 and 582 are first order filters. Filters 572 15 and 582 may be realized, for example, with capacitors and resistors. An added benefit is that high pass filters provide DC isolation and some rejection of the voiceband signal.

Figure 6 illustrates one embodiment of hybrid 260 in greater detail. Hybrid network 620 includes a hybrid input port, a hybrid output port, and a 20 receive path input port. Driver 612 is coupled to provide the upstream data signal to the subscriber line through resistors R_D . The driver is capacitively coupled to the hybrid input for provide the upstream data signal to the hybrid network. The receive path is capacitively coupled to the subscriber line for receiving the composite signal including the upstream and downstream data

signals. Capacitors C_{41} , C_{42} , C_{44} , and C_{45} provide DC isolation for the hybrid's transmit and receive paths.

A first order model approximates the subscriber line impedance, $Z(s)$, as a series-coupled resistor R_x and capacitor C_x . The transfer function from

5 the driver 610 output to the receiver input has a dominant zero and a dominant pole as follows:

$$h_{drv_to_rxin}(s) = \frac{Z(s)}{R_D + Z(s)} = \frac{1 + sC_xR_x}{1 + sC_x(R_x + 2R_D)}$$

The transfer function between the hybrid receive and hybrid output is as follows:

$$10 \quad h_{rxin_to_HYBout}(s) = K_{rx} \cdot \frac{s}{s + HYB0}$$

such that the composite transfer function between the hybrid receive input and the hybrid output is a high pass function as follows:

$$h_{drv_to_rxin_toHYBout}(s) = K_{rx} \frac{1 + sC_xR_x}{1 + sC_x(R_x + 2R_D)} \cdot \frac{s}{s + HYB0},$$

where K_{rx} is the receive path gain (i.e., the hybrid circuit gain) and HYB0 is the pole of the high pass transfer function (i.e., the hybrid zero). HYB0 is used to minimize the effect of the dominant zero ($1 + sC_xR_x$) due to line impedance. HYB0 is below 140 KHz so the receive signal is not significantly affected by the hybrid zero. The receive path gain, K_{rx} , is controlled by the values of capacitors CR and CF as follows:

$$20 \quad K_{rx} = \frac{CR}{CF}$$

A hybrid path is used to minimize the effect of the remaining dominant pole ($1 + sC_x(R_x + 2R_D)$). The transfer function from the driver output to the hybrid output through the hybrid path is as follows:

$$h_{drv_to_HYBin_to_HYBout}(s) = K_{HYB} \cdot \frac{s}{s + HYBP}$$

HYBP is the hybrid pole. K_{HYB} may be adjusted for different receive path gains (K_{rx}) such that

$$K_{HYB} \cdot \frac{s}{s + HYBP} - K_{rx} \cdot \frac{1 + sC_xR_x}{1 + sC_x(R_x + 2R_D)} \cdot \frac{s}{s + HYB0} \rightarrow 0$$

The hybrid path gain, K_{HYB} , is controlled by the values of capacitors CG and

5 CF as follows:

$$K_{HYB} = \frac{CG}{CF}$$

The hybrid pole and zero are determined from the components of Figure 6 as follows:

$$HYB0 = \frac{C_{41} + C_{45}}{R_Z \cdot C_{41} \cdot C_{45}}$$

$$HYBP = \frac{C_{42} + C_{44}}{R_P \cdot C_{42} \cdot C_{44}}$$

Thus

$$K_{HYB} = K_{rx} \cdot \frac{1 + sC_xR_x}{1 + sC_x(R_x + 2R_D)} \cdot \frac{s + HYBP}{s + HYB0}$$

Substitution indicates K_{HYB} may be calculated as follows:

$$K_{HYB} = K_{rx} \cdot \frac{1 + sC_xR_x}{1 + sC_x(R_x + 2R_D)} \cdot \frac{s + \frac{C_{42} + C_{44}}{R_P \cdot C_{42} \cdot C_{44}}}{s + \frac{C_{41} + C_{45}}{R_Z \cdot C_{41} \cdot C_{45}}}$$

15 Hybrid 260 is thus a tunable low order hybrid. HYB0 is adjusted to cancel or ameliorate the zero defined by $(1 + sC_xR_x)$. When properly tuned, hybrid 260 is effectively a first order hybrid. HYBP and K_{HYB} are adjusted to cancel the effects of gain K_{rx} and pole $1 + sC_x(R_x + 2R_D)$. The optimal K_{HYB} is programmed accordingly. Circuitry 620 is a capacitor coupled amplifier which
20 has almost no noise from the passive components. The feedback resistors (RF) for 622 serve as DC stabilizing resistors. When the values for RF are large, the resistors will contribute little noise. The first order hybrid may be implemented on the same substrate as the remainder of the analog front end.

The high pass filter 414 following the hybrid further reduces the remaining driver signal before the receive signal is digitized by ADC 420.

A crystal oscillator is used to generate all the necessary clock signals for the digital portion of the analog front end. Although there is some choice in selection of the crystal frequency, the selected crystal must support the DMT 4.3125 KHz channel separation. Accordingly, the selected crystal should have a resonant frequency that is a multiple of 4.3125 KHz. In one embodiment, the selected crystal is a 35.328 MHz crystal.

Referring to Figure 2, analog front end 218 can be configured to support

10 a voiceband channel for voiceband modem communications concurrent with the non-voiceband (i.e., xDSL) communications. The voiceband channel circuitry may be integrated on the same integrated circuit die as the xDSL circuitry, if a single clock is used for synchronization. Voiceband communications, however, have a bandwidth of approximately 4 KHz and
15 nominal baseband sampling rates of 8 KHz. These values are not immediately compatible with the requirements of the xDSL circuitry having 4.3125 KHz bandwidth subchannels. The use of multiple independent clocks may otherwise interfere with the operation of at least one of the xDSL or voiceband circuitry due to intermodulation and other undesirable effects
20 resulting from multiple clocks.

Referring to Figure 2, sample rate conversion is used to derive the other clock frequencies so that multiple independent clocks can be avoided.

Synchronous sample rate conversion is used to achieve the various clock rates through combinations of upsampling by an integer factor and

25 downsampling by an integer factor. Sample rate conversion block 210

includes a transmit path and a receive path which are illustrated in greater detail in Figures 7-8. Actual sample rate conversion is handled by blocks 741, 750, and 752 of Figure 7 and blocks 844, 850, and 860 of Figure 8. The remaining blocks condition the signal before or after sample rate conversion.

5 In one embodiment, the voiceband portion of the analog front end supports a number of clock frequencies including those in the set $S=\{7200, 8000, 8229, 8400, 9000, 9600, \text{ and } 10,286 \text{ Hz}\}$. A variety of clock frequencies allows flexibility in negotiating a common baud rate with another modem. A universal conversion speed (i.e., intermediate frequency) is selected which

10 can be derived from any of the required original rates by integer upsampling and downsampling. The appropriate upsampling and downsampling integers may be programmatically selected depending upon what baud rate the user is attempting to communicate with. Subsequent digital processing is handled at pre-determined intermediate frequencies independently of the

15 original frequency.

The sampling frequency f_s , is selected such that $f_s = \frac{f_o}{n}$, where n is an integer and f_o is divisible by n such that $f_o \bmod(n) = 0$. In one embodiment, $n=32$ so that $f_s = \frac{35.328\text{MHz}}{32} = 1.104\text{MHz}$. The conversion speed, f_{audio} , is selected judiciously such that for every f_i in S there exists integers M_i and N_i such that $f_i = \frac{M_i}{N_i} f_{audio}$

20 In one embodiment, the selected universal conversion speed is 9600. Table 900 of Figure 9 illustrates nominal M and N values when converting analog signal rates to the same universal conversion speed of 9600 for further digital processing. Common multiples of M and N help to reduce the

25 number of distinct Ms or Ns while maintaining the same ratio. Thus instead

of the fully reduced values of column 2, the use of reducible ratios requires only 3 distinct M: 6, 14, and 16. Although a least common multiple could be used to further reduce M, higher values may not be desirable for N or for digital filters such as finite impulse response filter 750. Thus column 3
5 illustrates the ratios optimized to minimize the number of distinct Ms for small M. The ratios provide precise conversions except for 10,286 and 8229 which have errors less than 0.5 Hz as required by ITU standards.

Referring to the transmit path 700 of Figure 7, the optimum M and N are selected from column 3 in accordance with the selected analog speed from
10 Table 900. Thus, for example, if the selected analog speed is 7200, then M = 16 and N = 12. Block 752 upsamples the audio signal by M. The upsampled signal is provided to finite impulse response (FIR) filter 750. The reduction in the total number of Ms reduces the number of filters required and therefore the memory required for filter coefficients. After digital filtering, the filtered
15 signal is downsampled by N in block 742. The data is now at the universal conversion speed of 9600. The downsampled signal is passed to compensator 740 to shape the frequency response before upsampling. Compensator 740 provides frequency compensation at high frequencies.

In block 732, the compensated signal is upsampled by 5 and then
20 filtered by infinite impulse response filter 730 to eliminate images introduced by the upsampling. The filtered signal is upsampled again by 23 in block 722. The upsampled signal is then passed to interpolator 720 to eliminate the images introduced by the second round of upsampling. In one embodiment, the 23 times upsampling and the interpolation are performed by the same
25 filter. The interpolated signal is then provided to DAC 710.

The frequency corner for the interpolator is selected to ensure adequate attenuation of signal images as well as no aliasing in the decimation process. This requires the FIR 750 cut-off frequency to be the smaller of the interpolation and decimation cut-off frequencies.

5 Figure 10 illustrates the operation of FIR 750. The FIR comprises a plurality of samples 1000 of the input data. Each sample 1000 is delayed with respect to its adjacent sample. Each tap has an associated coefficient 1010. The products of the coefficients 1010 and their associated tap values (from a preceding sample) are computed and then summed by adder 1020 to produce
10 a result. In the particular instance, every M coefficients 1010 is used to calculate a new output value. The input data shifts N places between calculation of output values. When interpolating, a number of the input sample values will have a zero value. Figure 11 illustrates a polyphase pointer and serial interface pulse generator 1100 for the FIR of the AFE's
15 transmit path.

Figure 8 illustrates the receive or downstream path 800. ADC 810 samples at 1.104 MHz. The resulting digital signal is passed to decimator 820. The decimated signal is downsampled by 23 in block 822. In one embodiment, the functions of downsampling by a factor of 23 and decimation
20 are performed by the same filter. At this point the data is being communicated at 48 KHz. Block 830 scales the data values, if necessary, to accommodate computational requirements of filters 840 and 850. In one embodiment, block 830 scales the data values by 64. Infinite impulse response filter (IIR) 840 is applied to remove the outband signals that may be aliased to

the baseband range of 9600 Hz to 48 KHz. The resulting data is downsampled by a factor of 5 in block 842. At this point the data is a 9600 Hz signal.

M and N are selected again based on the pre-determined analog speed from Table 900. Block 844 upsamples the signal by N before providing the 5 upsampled signal to FIR filter 850. The signal is then downsampled by M so that it has the proper rate for the voiceband channel. Figure 12 illustrates one embodiment of a polyphase pointer for the receive path of Figure 8.

In the preceding detailed description, the invention is described with reference to specific exemplary embodiments thereof. Various modifications 10 and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.